



(72) HENRI, Christopher J., CA

(72) COX, Robert D., CA

(72) RUBIN, Richard K., CA

(71) HENRI, Christopher J., CA

(71) COX, Robert D., CA

(71) RUBIN, Richard K., CA

(51) Int.Cl.⁶ G09B 5/02, G09B 19/24

(54) **ACCES A DES FICHES D'ENSEIGNEMENT BASEE SUR LE
WEB DANS UN ENVIRONNEMENT DE RADIOLOGIE SANS
FILM**

(54) **WEB-BASED ACCESS TO TEACHING FILES IN A FILMLESS
RADIOLOGY ENVIRONMENT**

ABSTRACT

This paper describes the design and implementation of a low-cost PACS conforming to the DICOM 3.0 standard. The goal was to provide an efficient image archival and management solution on a heterogeneous hospital network as a basis for filmless radiology. The system follows a distributed, client/server model and was implemented at a fraction of the cost of a commercial PACS. It provides reliable archiving on recordable CD and allows access to digital images throughout the hospital and on the Internet.

Dedicated servers have been designed for short-term storage, CD-based archival, data retrieval and remote data access or teleradiology. The short-term storage devices provide DICOM storage and query/retrieve services to scanners and workstations and approximately twelve weeks of 'on-line' image data. The CD-based archival and data retrieval processes are fully automated with the exception of CD loading and unloading. The system employs lossless compression on both short- and long-term storage devices. All servers communicate via the DICOM protocol in conjunction with both local and 'master' SQL patient databases. Records are transferred from the local to the master database independently, ensuring that storage devices will still function if the master database server cannot be reached.

The system features rules-based work-flow management and WWW servers to provide multi-platform remote data access. The WWW server system is distributed on the storage, retrieval and teleradiology servers allowing viewing of locally stored image data directly in a WWW browser without the need for data transfer to a central WWW server. An independent system monitors disk usage, processes, network and CPU load on each server and reports errors to the image management team via email.

The PACS was implemented using a combination of off-the-shelf hardware, freely available software and applications developed in-house. The system has enabled filmless operation in CT, MR and ultrasound within the radiology department and throughout the hospital. The use of WWW technology has enabled the development of an intuitive web-based teleradiology and image management solution that provides complete access to image data.

Keywords: PACS, DICOM, recordable CD, image management, filmless radiology

ABSTRACT

This paper describes the incorporation of radiology teaching files within our existing filmless radiology Picture Archiving and Communications System (PACS). The creation of teaching files employs an intuitive World Wide Web (WWW) application that relieves the creator of the technical details involving the underlying PACS and obviates the need for knowledge of Internet publishing. Currently, our PACS supports filmless operation of CT, MRI, and ultrasound modalities, conforming to the Digital Imaging and Communications in Medicine (DICOM) and Health Level 7 (HL7) standards. Web-based teaching files are one module in a suite of WWW tools, developed in-house, for platform independent management of radiology data. The WWW browser tools act as liaison between inexpensive desktop PCs and the DICOM PACS. The creation of a teaching file is made as efficient as possible by allowing the creator to select the images and prepare the text within a single application, while finding and reviewing existing teaching files is simplified with a flexible, multi-criteria searching tool. This efficient and easy-to-use interface is largely responsible for the development of a database, currently containing over 400 teaching files, that has been generated in a short period of time.

Keywords: teaching files, world wide web, filmless radiology, PACS

1. INTRODUCTION AND BACKGROUND

The Department of Radiology at the Montreal General Hospital (MGH) provides radiological services for our 500-bed, level-1 trauma centre, as well as for external clinics in neighbouring hospitals such as the Royal Victoria Hospital and the Montreal Jewish General Hospital. Our department has been in filmless operation since February of 1997 for CT and MR, and since October of 1995 for ultrasound. Since then, the department has performed approximately 72,000 filmless exams in these modalities. To accomplish this, a modular PACS was developed in-house to provide connectivity between scanners, commercial workstations, and storage servers for on-line and off-line data. As well, a recent development has allowed us to interface our PACS with the Hospital Information System (HIS) and we are currently developing our own Radiology Information Service (RIS) to provide electronic reporting functionality conforming to HL7 standards. A WWW server provides access to our PACS for computers with Internet connectivity. This allows non-DICOM clients (such as inexpensive PCs) to view images and reports, issue DICOM move commands, print to paper, and issue pre-fetching instructions to the system.

7

The distributed DICOM storage servers are based upon inexpensive Intel-Pentium platforms running a public domain variant of UNIX, called LINUX. DICOM conformance is achieved by developing tools employing the Mallinckrodt Institute of Radiology Central Test Node (CTN) library, and a public domain database (Mini SQL, Hughes Technologies) is used to store patient and image information. We currently employ four such servers (called modality servers) for three CT scanners, one MR imaging scanner, and nine ultrasound acquisition stations; while two servers provide archiving and off-line retrieval functions respectively.[1]

Only modality servers receive images directly from the scanners (i.e., scanners never send images directly to workstations), since only modality servers update a master database which tracks the location(s) of every image in the system along with all its relevant information. These modality servers are equipped with five 9 GB disks in a level-5 RAID array to serve as short term on-line storage facilities (typically 12 weeks), which support the DICOM Storage and Query/Retrieve service classes as both provider (SCP) and user (SCU).[2] As well, they automatically route images to specific workstations based on modality and anatomy, which is how our radiologists are organised. Additionally, a copy of the image is routed to the archival server to be placed on compact disc recordable (CD-R) media for long term storage.

¹ Further author information -

R.K.R.(correspondence): Email: rrubin@rad.mgh.mcgill.ca; Telephone: 514-937-6011 ext. 3386; Fax: 514-934-8229



High-end reporting workstations are accessible within the radiology department and high-end reviewing stations are situated in key locations around the hospital (e.g., operating and emergency rooms, clinics, and certain conference rooms). However, a more economical alternative was needed for those clients who previously had access to film but could not benefit from these workstations. Most of the computing infrastructure of our hospital, as well as those of client hospitals and external clinics, are comprised of low-end personal computers. Due to the rapid proliferation of the Internet and the minimal requirements for PCs to run a web browser, a web-based solution was adopted to satisfy these users. Therefore, WWW access to the PACS addresses the needs of a diversified and distributed user group by leveraging the popularity and familiarity of browser interface technology.

A teleradiology server provides the web-based access to the PACS using similar architecture as the modality servers, including the Apache HTTP server. In our initial design, all web communication occurred between a browser and the teleradiology server.^[3] Therefore all images from modality servers needed to be routed to the teleradiology server and converted to JPEG format for in-line viewing. The result was a bottleneck in an otherwise distributed system. A recent innovation to WWW access provides distributed in-line image viewing. Now, a web browser will be automatically re-directed to the appropriate modality server that stores the images in full resolution DICOM format as well as the compressed JPEG format for web-based viewing. The CPU requirement to convert the images is now distributed as well, resulting in a low latency time for image conversion. Nothing changes from the point of view of the user, who still connects to a single Internet address, except there is much less waiting because much more data is stored in pre-converted JPEG format.²

2. IMPLEMENTATION

Our teaching file server became fully operational in September 1997. Collaboration between the PACS developers and radiologists of our department resulted in a tool to easily identify specific image series, add descriptive text to them, and have a fully searchable index of them. The result is global Internet access to the teaching files for both creator and audience, since all interaction is provided by a web server.

Teaching file data and web access to the teaching files are localised to one server (also an Intel-Pentium platform running LINUX). The server is a DICOM storage SCP as well as a WWW server. Images moved from a modality server to the teaching file server are stored in DICOM format and then converted to JPEG format. The teaching file server behaves differently than a modality server since patient names are blanked in the header of each image and no patient names are stored in its local database, nor does it update the master database to indicate which images it is storing. DICOM files are kept so that a user may adjust the window and level settings of the in-line images. As well, the DICOM formatted images may also be downloaded to the local hard disk of the client, via the web-browser, where more powerful viewing software may be used.

Common Gateway Interface (CGI) scripts written in Perl, provide the tools to create, modify, search, and view the teaching files. As well they control access privileges to restricted functions. For instance, only the owner of a teaching file, or accounts with ownership rights may make modifications to a file or delete it and a guest account may not create a teaching file. One particularly useful tool, restricted to department members, is the ability to compile a group of teaching files and have them automatically saved on a CD-R. This allows visiting radiologists to bring their cases with them when they leave the department, since film is not otherwise available. This also enables radiologists to prepare rounds for facilities not equipped with a network connection.

3. DESCRIPTION

The teaching file server was designed to fit in as part of a larger toolkit to allow web-based clients to interact with our filmless PACS. In so doing, re-training was not necessary except for those functions unique to creating teaching files. As well, even though the teaching file server is a physically separate entity, the way in which a user interacts with the web tools appears seamless. The interactions specific to teaching files can be grouped into three categories: 1-teaching file creation, 2-teaching file search and review, and 3-teaching file maintenance and utility.

² For a detailed discussion of our PACS design, please see our companion paper in these proceedings by R. D. Cox, *et. al.*, titled "DICOM-compliant PACS with CD-based image archival."

3.1. Teaching file creation

The first step in creating a teaching a file is to identify the series of images that will be grouped within one file. Radiologists are also able to add scrapbooks (single series collections of clinically significant annotated images) to teaching files, which are created on the review workstations and then pushed onto a modality server, which treats them as any other new series it receives (i.e., updates the master database, creates JPEG images and sends the files to be archived). To identify and add any series to a teaching file, the creator first logs on to the teleradiology web server using a user name and password. A search utility allows the appropriate series to be found by descending a hierarchical listing of patients, studies, and then series. With a list of series shown, the user selects one or more series and chooses 'Create Teaching File' from a pull-down list of possible actions. If the user account did not possess the right to create a teaching file, this option would not be present. As shown below in figure 1, the user has selected three scrapbooks¹. Just below the selection window, is a pull-down menu of actions to perform on the selected series. Here, 'Create Teaching File' is chosen, and when 'Submit' is clicked, the user will be taken to a teaching file input form where textual and other information may be entered.

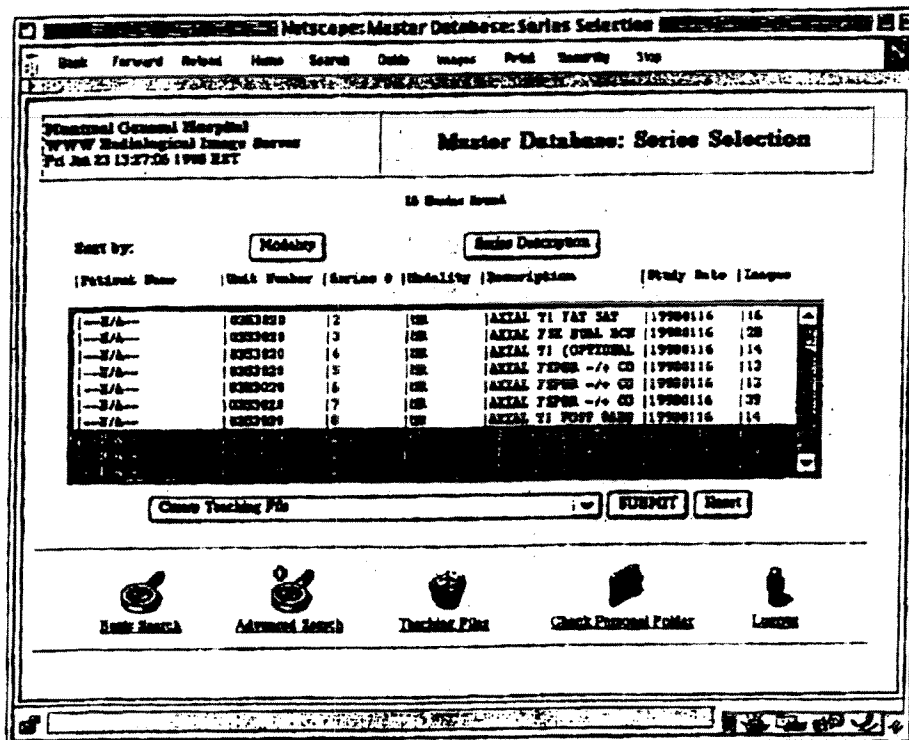


Figure 1: Series selection web page. Shown here, three scrapbooks are selected and are about to be grouped into one teaching file.

A check is done to see if any teaching files already exists with the same unit number (medical record number) as any of the series being added. If one or several matches are found, a list is displayed and the option is presented to append the new series to one of the existing teaching files, or to create a new file (figure 2).

If the option to append the new series to an existing file is chosen, the form is filled out automatically with the information from the teaching file. This is useful for progressively presenting a study. The input form consists of text input fields, pull-down menus and check-box fields. Patient information such as age, sex and unit number, as well as the information about the images (modality, anatomy, name of physician) is automatically copied as well and shown at the top of the page.

¹ Scrapbook modality is denoted as 'Other' (OT)

The creator of the teaching file has twelve fields to fill out, but only three are text inputs. The text input fields are for entering descriptions of clinical findings, descriptions of image findings (figure 3), and a list of keywords delimited with the slash character ('/'). Keywords can also be selected from a list of previously entered keywords, and that list is updated whenever a new keyword is adopted. The text inputs place no restriction on the users concerning length or content. A pull-down menu with a list of exams is provided. For example, a MR abdominal case would have options such as: dermoid, MRA, spleen, etc.. If the radiological study is proven or unproven or if it needs follow-up, this can be indicated with a check-box.

Teaching File(s) Already Exist For Selected Unit Number(s)

UNITID PACS Code Unit Number Keywords

1548	1	0001430	Liver Lesion
------	---	---------	--------------

Modify Selected Teaching File Submit

Figure 2: Existing teaching files found containing the unit number of the new series. The user may now create a new file, or append the images to an existing one.

If the creator is not a staff radiologist, but rather is a resident or a fellow, a pull-down menu of staff radiologists can be used to link one of them to the file, giving the staff member ownership rights. This list of staff radiologists is obtained from the main user database, and no changes to the scripts are needed when new accounts are added or removed. Two text fields are provided for ACR coding but, so far, not a single teaching file creator has entered information here. The consensus is that keywords, and the ability to search for them, are more meaningful and, hence, more useful to our radiologists.

Image Quality: ☐ 1 ☐ 2 ☒ 3 ☐ 4 ☐ 5 (Best)

Description Of Image Findings:

The US demonstrates a hyperechoic lesion in segment 4, typical for hemangioma. The second lesion is involving segments 5,6,7, & 8. It is hyperechoic and has a very prominent central hyperechoic nodule. No obvious discontinuous diaphragm sign. Could this be fat in a lesion, or blood?
MR confirms the segment 4 hemangioma, with typical signal intensities (low on T1,

Figure 3: Portion of teaching file creation form.

Study interest and image quality can be ranked from 1 (lowest) to 5 (highest) by clicking the appropriate check box (figure 3) and the images can be flagged as being of print quality. This is important for our users wishing to search this facility to find interesting cases with high quality images for presentations and publications.

When the teaching file creator is satisfied with the form, a button at the bottom of the page submits the information to a script on the teaching file server. The receiving script checks that certain fields are filled in, and if key information is missing, notifies the user and returns them to the input page to enter the missing information. Otherwise, the script begins searching the master database for the location of the images in the PACS and then issues DICOM move commands to the teaching file server from all the necessary modality servers. If images are only stored in off-line archives, the archive retrieval server will be instructed to retrieve and send the required images to the teaching file server. When the images arrive at the teaching file server, they are stored in DICOM format (12-bit data), and a JPEG format image is also created and stored.

The data from the input form is added to a database on the teaching file server, or in the case where an existing file is being modified, a database update is performed. The relational database is structured into individual tables for the teaching file text, image information, and a table of keywords. For each teaching file, there is one teaching file text table entry with a unique identification number. This identification number links one or more table entries in the image information table. The keyword table is independent, and is merely used to track the list of the current keywords being used.

3.2. Teaching file search and review

Access is gained to the teaching file server through our login page on the teleradiology server. The users must identify themselves using a user name and password combination. Guests may use a generic account with few privileges beyond viewing teaching files, and patient names are withheld during any of the limited access to the rest of the system.

The first page shown upon entry is a menu of functions available to the user. The first menu item brings the user to a search facility shown in figure 4. This is presented as a fill-out form that allows users to build their queries. The form is designed to be as efficient as possible for entering queries ranging from simple to complex. The form can be logically divided into two parts. The top portion allows the user to build a sentence-like query by using pull-down menus and a text input field. The first menu contains a list of 15 text-searchable fields, with 'Keywords' at the top of the list and pre-selected since it is used most often. The next menu contains a list of logical statements: 'contains', 'begins with', 'is', 'is greater than', and 'is less than'. By default, the 'contains' option is selected. Following these two menus is a text entry field for the user to use to complete the sentence. Multiple items can be entered delimited by the '/' character, and the conjunction is a logical 'or'. At this point, the user may submit the query, but there is the option to limit the results further by using the next line below which is similar to the first sentence-like query. The two can be combined in a logical 'and' or a logical 'or' function.

Search Reset

Keywords contains liver

AND

Clinical Findings contains cyst

Specialty: ☐ All Cases ☒ Abdominal ☐ Bone ☐ Chest ☐ ENT ☐ Neuro ☐ Spine ☐ Other

Status: ☐ Proven ☐ Unproven ☐ Needs Followup

Image Quality: ☐ 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 (Best)

Study Interest: ☐ 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 (Most)

Good For Prints: ☐

Examination Type:

Search Reset

Figure 4: Teaching file search form. The resulting query of the form above will select all teaching files that have keywords containing 'liver' and clinical findings containing 'cyst' and belonging to the 'abdominal' specialty.

The second portion of the query form is composed of check-box type input fields. These fields act as 'and' functions with the sentence-like queries. Using the check-box inputs, the search can be limited to radiologist specialty (abdominal, bone, neurology, etc.), exam status, image quality, study interest and print quality. As well, a pull-down list of exam types is given.

The page following the query form presents a list of teaching files matching the search criteria (figure 5). At this point a teaching file can be viewed, modified, deleted or marked for addition to a CD-R (assuming the account has all the proper privileges). Guests may only view teaching files, and only owners may modify or delete a teaching file.

Sort by:

TF ID | ACR Code | Unit No | Keywords

128	857838		Liver atrophy/LHB/STUDIES/838
127	1127789		Amastomias liver transplant
104	0286197		Liver metastases
136	1145432		polypoid liver disease
321	1082486		Hydatid cyst, liver

Figure 5: List of teaching files matching search criteria.

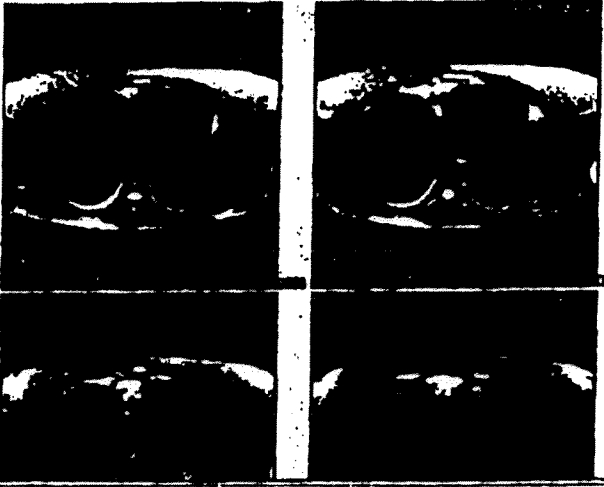
Choosing to view a teaching file brings up a web page containing all the information of the chosen file. At the top is a hyper-link called 'View Series'. This will open a new window displaying the images of the series in a tiled layout. Controls

MetaImage In-Run Viewing

Back Forward Refresh Home Search Help Images Print Summary Stop

Patient Name	Unit Number	Date Of Exam	Study Description	Series Description	Number Of Images	Referring Physician
-N/A-	1155-04	19980115	A LIVER	TYPE LIVER LESION	35	PAPADATOS

218218 (400)
Image Number
200 100



TF Creator:
Chakovsky

Radiologist:
Reichold

Examination Type:
Liver

Teaching File Description:
21 year old female with mild EOP cholecystitis. Large hepatic mass.

Status:
Proven

Image Quality:
5

Description Of Image:
Extensive heterogeneous mass of the RT and middle hepatic veins, and mass inferiorly. All vessels & biliary ducts well developed.

Keywords:
Liver Hepatoma

Study Interest:
5

Print Description:
Excellent example of Hepatoma.

Figure 6: Teaching file display. Shown are two overlapping windows, one containing the images of the teaching file and the other containing the description.

at the top of the page allow the user to modify the size of the images, turn image numbering on and off, modify the image sort order and adjust the brightness and contrast by entering new window and level values. When scout images are available.

a single image or a whole series can be referenced to its corresponding scout to produce a scout image with slice lines indicating the intersections of the images with the scout (figure 6).

All the programs for viewing images, re-sorting, adjusting brightness and contrast, and scout referencing were recycled from the in-line viewing application of our web-based PACS interface. Therefore, no extra programming was required for this interface. In general, the teaching file server was designed to look and behave much like the PACS interface to save development and training time. As an added benefit, many improvements or additions to the PACS application apply to the teaching file server as well.

3.3. Teaching file maintenance and utility

Managing teaching files is a task left to the staff radiologists, fellows and residents, with the exception of server specific issues. This includes content quality control and maintaining the list of keywords. To help keep keywords as consistent as possible, a keyword database editor application allows for the removal, modification and the ability to coalesce similar keywords into one. This not only changes the list stored in the keyword database, but alters the keyword entries in the teaching files as well. For instance, changing 'abd. mass' to 'abdominal mass' will remove 'abd. mass' from the database and add 'abdominal mass' unless it already exists. All teaching files with keyword entries matching 'abd. mass' will also be updated with the new keyword. This makes search results matching a single keyword more complete since multiple spellings and abbreviations are eliminated.

One very useful feature of our teaching file system is the CD-R creation facility. Over time, any user with the proper privileges can mark teaching files to be placed on a personal CD. Selecting one or more teaching files from the query result list and choosing 'Add Teaching File to my CD-R' does this. Yet another database table keeps track of the CD selections of each user, as well as the amount of space required to store the selected teaching file. Eventually, the CD will reach capacity and the user will be notified that no more files can be added to their current CD. The user may then choose to burn the CD as is or edit the contents using a simple interface. From the point of view of the user, having a CD burned requires just one button click from the main teaching file menu. This starts a script that looks up the contents of the CD and makes a copy of every image into a separate directory on disk. Then it makes one hypertext mark-up language (HTML) file for each selected teaching file containing all the teaching file information. Another HTML file is made for each image series. Hyperlinks are added to link the image HTML file with the teaching information HTML file. Finally, an index file is made, listing all the teaching files stored on disk with hyperlinks to each one. Unfortunately, this replaces the web-based search tool since it is not possible to add it to the CD because all searching is done on the server side of a web connection. Included on the disk is a copy of Netscape for both Macintosh and Windows 95, so the CD is all the user needs to view the files on a desktop PC running these systems.


Overnight, the CD image on the teaching file server is copied to a CD-R archive station and scheduled to be saved on a CD-R. By the following day, the CD-R is ready to be picked up.

4. RESULTS


Detailed statistics are kept by our web-based PACS application concerning actions performed by every user using the system. These statistics allow the system developers to monitor system activity to determine which features are most used, and therefore, may need performance improvements to keep pace with demand. Under-used functions could be explained in two ways. A function may not be well used because it simply is not needed, or it just may not function efficiently enough to be worth using. In any case, more than once, auditing usage trends have identified areas which need performance or interface improvements.⁴ In the event of any suspected misuse of the system, detailed audit trails are available for further investigation.

The teaching file related functions also add to the history of usage statistics kept for the web-based PACS application. Examination of these records shows that there now are 577 image series belonging to 485 teaching files on the server. 20 different users have created these since August of 1997. The median for new teaching files created per month is 53, and the median for the number of different creators per month is 7 (see graph 1). The break down of teaching files by speciality is 85% abdominal, 6% bone, 6% neurology, 2% spine, and 1% chest.

⁴ For a detailed discussion of the performance of our PACS, please see our companion paper in these proceedings by C. J. Henri, et al., titled "Acquisition and analysis of usage and data-flow statistics for a DICOM-compliant PACS."

New teaching files created per month**Graph 1: New teaching files created per month**

A total of 72 users have performed 1900 teaching file queries using the teaching file search utility. The median for the number of queries per month is 376, and the median for the number of different users performing those queries is 23 per month (see graph 2 below).

Teaching file queries by month**Graph 2: Teaching file queries by month****5. DISCUSSION**

The results for the creation of new teaching files show a peak of 213 in September and then tapered-off to a more constant level of around 50 new teaching files created per month. This is due to a batch conversion of teaching files stored in another database, which also explains the large skew of 85% abdominal speciality content. The previous database was only for MR abdominal cases, and did not store any images. So the immediate benefit of converting the older files to the new web-based form was the addition of images as well as a much broader potential audience (the previous database was only accessible

from Macintosh computers in our department). The creation of teaching files should remain somewhat constant for the near future, since residents must now create new teaching files to satisfy a quota that is part of their academic requirements.

Teaching file queries also show a steady trend of just under 400 queries per month (median is 376). August has few queries because the teaching files server had only just been created. User queries for January 1998 are below average despite the fact that this paper is being written only two thirds of the way through the month. This is probably due to the lack of activity in the department because of holidays and a major ice storm that caused the city and outlying regions to loose power for several days.

6. CONCLUSION

Our web-based teaching file server has three distinguishing features that make it a powerful and useful tool. The first is integrating the server with our existing DICOM PACS. There is no need to cut and paste images from one application to another. Instead, the creator can build teaching files using any series found in the master database, using a single, and consistent interface that is simple and easy to use. The second important feature is that the teaching file becomes immediately available after creation to anyone with Internet access. Most importantly, the creator does not have to know anything about creating web pages or any of the details concerning the location of images or the transfer of image series to the teaching file server. The third key feature of our teaching file server, is the fact that it was developed in-house along with our PACS. This allows the system to be tailored for the users by making modifications and introducing new features to keep pace with the ever-evolving requirements of the system.

ABSTRACT

This paper describes the design and implementation of a low-cost PACS conforming to the DICOM 3.0 standard. The goal was to provide an efficient image archival and management solution on a heterogeneous hospital network as a basis for filmless radiology. The system follows a distributed, client/server model and was implemented at a fraction of the cost of a commercial PACS. It provides reliable archiving on recordable CD and allows access to digital images throughout the hospital and on the Internet.

Dedicated servers have been designed for short-term storage, CD-based archival, data retrieval and remote data access or teleradiology. The short-term storage devices provide DICOM storage and query/retrieve services to scanners and workstations and approximately twelve weeks of 'on-line' image data. The CD-based archival and data retrieval processes are fully automated with the exception of CD loading and unloading. The system employs lossless compression on both short- and long-term storage devices. All servers communicate via the DICOM protocol in conjunction with both local and 'master' SQL patient databases. Records are transferred from the local to the master database independently, ensuring that storage devices will still function if the master database server cannot be reached.

The system features rules-based work-flow management and WWW servers to provide multi-platform remote data access. The WWW server system is distributed on the storage, retrieval and teleradiology servers allowing viewing of locally stored image data directly in a WWW browser without the need for data transfer to a central WWW server. An independent system monitors disk usage, processes, network and CPU load on each server and reports errors to the image management team via email.

The PACS was implemented using a combination of off-the-shelf hardware, freely available software and applications developed in-house. The system has enabled filmless operation in CT, MR and ultrasound within the radiology department and throughout the hospital. The use of WWW technology has enabled the development of an intuitive web-based teleradiology and image management solution that provides complete access to image data.

Keywords: PACS, DICOM, recordable CD, image management, filmless radiology

1. INTRODUCTION

The Montreal General Hospital is a 500 bed teaching facility that forms part of the McGill University Hospital Community. The hospital's Department of Diagnostic Radiology averages 160,000 examinations per year in conventional X-ray, angiography, mammography, MRI, CT and ultrasound.

Based on the success of a filmless ultrasound mini-PACS that was developed in-house,¹ a project was initiated to include MRI and CT in a distributed image management system which would allow the department to store conventional film in these modalities as well. While the ultrasound system consisted of Apple Macintosh and a Sun SPARCstation server connected by a local area network using the Appletalk protocol, the project was made to direct further PACS development towards supporting the DICOM 3.0 standard.

Other author information: (Send correspondence to R.D.C.):
Dept. of Diagnostic Radiology
Montreal General Hospital
1630 Cedar Avenue
Montreal, Quebec H3G 1A4
Canada
E-mail: rcox@rad.mgh.mcgill.ca

2. SYSTEM OVERVIEW

2.1. Distributed versus Centralized

From the outset, we chose to employ a distributed client/server design in our PACS development. This design provides several key advantages:

- Data are distributed throughout the network, and with proper design, the image data are already located where they will be required.
- There is no requirement for a 'super-server' to handle the entire PACS load, as both network and computational load are distributed among several servers.
- Servers can employ cheaper hardware as they do not need to be as powerful as a single centralized server.
- A distributed design can be easily scaled to suit growing demands by simply installing additional servers or expanding network segments to meet the need.
- The PACS can be designed with redundancy of both data and hardware, reducing or eliminating single points of failure.

A major drawback of distributed PACS is the complexity of the design and implementation. Image data can be found in a variety of locations and the applications that access them must be able to choose the optimal source for data they wish to transfer or process. This has increased the complexity of the developed system. Fortunately, this increase has remained transparent to the user community who see the PACS as a single entity. In essence, the goal is to design a distributed system with all its inherent advantages, but have it appear as a centralized system from the perspective of the user.

While data storage and access are distributed, it was necessary to design a central database to contain patient summary and image location information of all data scanned into the PACS network. This was done to relieve the user from having to know where in the system the image data reside. As image data are acquired and transferred from the scanner to a dedicated short-term storage device, or *modality server*, a summary of the given patient demographic, study and series information, and data location are transferred to this central database on the *master database server*.

2.2. System Components

The PACS is based on UNIX servers and makes extensive use of key UNIX features such as stable multi-processing and memory management, the cron facility for launching background processes, and a variety of public-domain servers, interpreters and development tools.

Figure 1 shows the features supported by each of the components of the PACS and the various interactions (DICOM, database and web-based) between them.

2.2.1. WWW Server

The primary user of the central database, is a full-featured set of web-based applications called the *PACS browser*.² With a web browser, a user can locate image data anywhere in the PACS network using a sophisticated search engine. Once located, images can be viewed within the browser, transferred to an alternate location such as a reporting workstation, printed on plain paper or downloaded at full resolution to a PC for off-line analysis. All data are transferred from source to destination transparently; that is, without the user needing to know where the data currently reside. The sole exception to this is the case where the data can only be found in long-term storage. Under these circumstances, the user is warned that the images will have to be retrieved from CD, as this procedure could take several minutes.

The PACS browser has gained wide acceptance and is used by every segment of the user community:

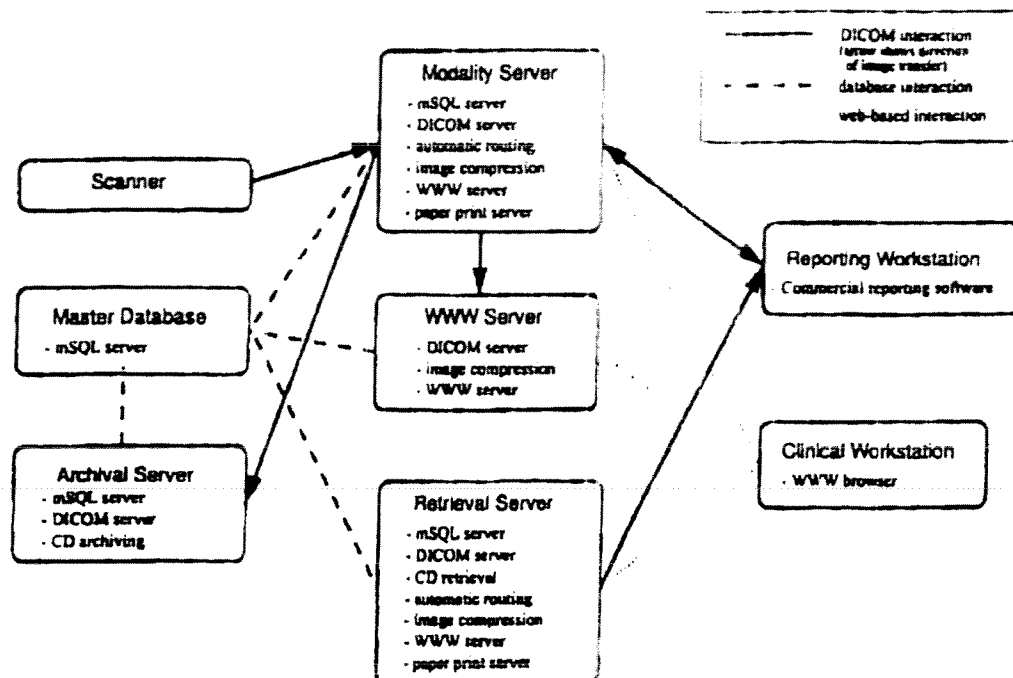


Figure 1. Shows features of various PACS components and interactions between them.

- Radiologists use the system to manage image data on the PACS, review work lists of cases in their specialty, engage in physician consultation, perform on-call review over a modem connection (teleradiology), and create teaching files* of interesting cases.
- Residents use the system as a learning tool, while gaining much needed exposure to a filmless radiology environment.
- Technicians, especially in ultrasound, use the system to look at a patient's previous studies before proceeding with the current scan.
- Filing clerks use the system to print images for patients on paper, prefetch the previous studies of patients for scans scheduled the following day, and fulfill requests pertaining to in-patient and emergency cases.
- Clinicians use the system for clinical review of cases, for patient consultation, and to provide their patients printed hard copies of image data on dedicated plain paper printers located throughout the hospital. Clinical review stations are used by surgeons in the operating room in place of film and a light box.
- The image management team uses the system to troubleshoot image storage and database problems, to manage system user accounts, and to check the contents of databases

2.2.2. The Generic Server

At the heart of the PACS is a generic DICOM-conformant server based on the Linux operating system (kernels 2.0.27-2.0.30, distributions RedHat 4.0-4.2). These servers run a suite of applications configured depending upon their role in the PACS. For example, a modality server runs an mSQL database server[†], a DICOM server based on

*See the companion paper, R.K. Rubin, et.al., Web-based access to teaching files in a filmless radiology environment, in these proceedings

†Hughes Technologies, Pty., Australia

the MIR CTN library,³ an automatic routing daemon, a custom C-MOVE process, a paper print server, a DICOM to-JPEG conversion process, and a WWW server to support the PACS browser. When a server is configured as a retrieval server, used to retrieve data from CD, it runs three additional processes: two graphical applications to service archived image transfer and print requests, and a daemon to process prefetches. Each server runs one or more scripts, written in Perl, to monitor disk usage and automatically delete data and update databases. All servers run a suite of system-monitoring scripts called Big Brother⁴ that track running processes, disk usage, processor load, swap space usage and several other parameters. Image management team members are notified by email if any of these parameters exceed tolerance.

The hardware configuration of a typical server consists of an Intel Pentium Pro 200 or P-II 266 MHz processor with an ultra-wide SCSI Redundant Array of Inexpensive Disks (RAID) controller and a fast Ethernet network interface card. Hardware and storage are configured to suit the needs of the application. For example, a modality server has five 9 gigabyte (GB) disks configured in a RAID level 5 array yielding about 32 GB of formatted storage capacity. Database servers employ two disks in a RAID-1 configuration which provides data redundancy without the write performance degradation exhibited by RAID-5. A retrieval server has a two-disk RAID-0 array which acts as a fast disk cache for data retrieved from CD-ROM.

2.2.3. Modality Server

Modality servers receive images from scanners and provide short-term storage within the PACS. There are currently two servers dedicated to CT, one to MR and one to ultrasound. Each machine is equipped with a RAID controller supporting a 5-disk RAID-5 array which provides about 32 GB of formatted storage. Each modality server also has an IDE boot disk, a fast Ethernet card and 192 MB of EDO RAM. The DICOM application is both a storage Service Class Provider (SCP) and Service Class User (SCU), as well as providing query/retrieve services.

Incoming images are stored on the RAID in a nested directory structure consisting of Patient ID and Study Instance UID. Individual image names consist of the series number, image number and a random number to prevent the inadvertent overwriting of data by the system. A local database is written with demographic, study, series and image information as well as the specific path to the image data.

The DICOM server application supports several options that are determined at run-time:

- Automatic Routing: images can be automatically routed to one or more destinations based on work-flow rules
- Automatic Archival: series can be tagged for automatic archival to a dedicated archival server.
- Immediate or delayed compression: images may be compressed immediately or have compression performed at a later time by a separate process.
- Teleradiology Functions: images can be tagged for conversion to JPEG format for in-line viewing via the local PACS browser.
- Automatic Prefetch: an incoming study can automatically trigger the transfer of a previous study from long term storage.

Modality servers also query the master database server to determine if stored records have already been transferred. Each modality server contains an exact copy of the master database structure. As images are received, the DICOM server writes appropriate records to the local master database. A separate process determines if these records are already contained in the remote master database, inserts or updates them as necessary and deletes the local copy. By decoupling remote database storage from the DICOM software, the modality server can continue to operate and store data even if the master database server is unavailable.

The disk usage of the modality server is monitored by a script that runs in two modes. During the day, the script monitors disk capacity and intervenes to delete data only to keep the disk below the maximum capacity threshold. At night, the same script deletes the oldest data from the RAID until it reaches the minimum capacity threshold which is calculated to give enough headroom for the scans of the following day. In this way, most of the overhead of data deletion and the associated database queries do not occur during peak hours.

2.2.4. Master Database Server

The master database server runs the mSQL server application and houses the large central database that describes all data scanned since the PACS came into use. This database is queried chiefly by applications of the PACS browser as it has no DICOM capability. It employs a RAID-1 array for fast data transfers and data redundancy. It receives updates from applications on the modality servers that transfer master database records.

The conceptual unit of data in the system is the image series. Each series in the PACS, as identified by a DICOM unique identifier or UID, has associated with it one or more location records in the master database. These location records tell PACS applications where the data can be found; e.g., on a modality server or on a specific volume of archival media. Applications can then choose the optimal location from which to retrieve the required data.

The master database server also houses an HL7 database which will become the core of the reporting system currently in development. The project encompasses the integration of mainframe demographic information and orders with radiology reports and scanned images.

2.2.5. Archival Server

The archival server is used to receive images for long-term storage and tape backup. The nature of the chosen archive medium, recordable CD, requires that data be written to each CD in a batch-mode recording process. Recording of CDs also requires the full use of the recording computer as any interruption of the data stream during the recording process will corrupt the CD, making it unreadable. These limitations led to a system design where CDs are recorded during the day and data are received from modality servers overnight.

The archival server is similar in many respects to a standard modality server without capabilities such as automatic routing, WWW in-line viewing or printing. It acts as a storage SCP to the modality servers which send all series tagged during the day for overnight archival. The incoming images are stored in directories corresponding to the CD number where they will be located, and are compressed as they come in. An independent script, run by the cron facility, monitors each active CD directory (there is one active directory for each modality, CT, MR and US). When a directory reaches the limit for a recordable CD (about 650 MB) the CD number is incremented, a new directory created, and the control database changed to reflect the new storage location for that modality.

After all modality servers have transferred their available data, another process runs which creates a table-of-contents file for the CD data and builds an ISO-9660 image file of the directory suitable for recording on CD. The spooled, compressed data is also written to a backup tape which can be used to rebuild a CD, if lost or damaged. An interactive program, run during the day, is used to record prepared image files onto blank CDs. This is the only part of the CD recording process that requires user intervention. All other operations are completely automated.

The system has been designed in such a way that the choice of archive media is transparent to the rest of the system. In order to change media, we would simply have to make minor changes to the archival server software to accommodate the recording method. As the PACS grows and the amount of stored data increases, the burden on existing archival and retrieval resources has become apparent. We are anticipating the move to alternate media such as Digital Versatile Disk (DVD) in the near future and have also been evaluating Digital Linear Tape (DLT) for near-line storage.

2.2.6. Retrieval Server

The retrieval server is used by filing clerks to fulfill requests for data to be retrieved from long-term CD archive. The retrieval server can be thought of as a fully equipped modality server where the storage SCU is an application that reads data from CD, rather than a scanner. Two graphical X-windows applications are used to retrieve and distribute (or print) images from CD. Several other applications run behind the scenes to schedule and process requests that can be carried out without user intervention. Once data have been stored on the disk cache of the retrieval server, they are converted to JPEG format to become available for in-line viewing via the PACS browser.

The retrieval cache is monitored by a script that deletes the least recently retrieved series when the cache reaches capacity. As with the modality server, the disk monitor runs in two modes, daytime emergency monitoring and nighttime disk maintenance to reduce the computational and disk I/O burden on the server.

2.3. PACS Features

2.3.1. Automatic Routing

To provide work-flow management to the PACS, an application was developed to automatically route images to the locations where they are required. As images are received from the scanner by the modality server, they are optionally tagged for automatic routing by placing an entry in a routing database. A daemon process monitors this database and performs all image routing according to a set of routing rules. In our department, the radiologists are specialized by anatomy, so reporting workstations have been assigned to anatomical categories of abdomen, bone, chest, ENT, neuro, and spine. As well, some workstations are shared between groups. The routing rules accommodate these work schedules by routing specific scans to specific workstations at specified times.

To accomplish this, scans are coded by anatomy at the scanner using the DICOM *Study Description* field. This information along with the date and time scanned and the workstation list are used to identify the routing destination. The images are then scheduled to be sent to the destination at the appropriate time as defined in the routing rules.

2.3.2. Custom C-MOVE Implementation

Two problems were noted with the generic DICOM C-MOVE service class provided by the CTN library. Firstly the DICOM server is a multi-process application so a duplicate C-MOVE request could be issued repeatedly to the server and each request would be carried out regardless of the fact the the same data had been previously requested. It was noticed early on that, because of the nature of the workstation interface, users were repeatedly transferring the same data, which taxed the server, the workstation, and the network. Secondly, the generic implementation of C-MOVE did not give adequate feedback to the workstations or PACS browser of the success or failure of the C-MOVE request.

To solve both these problems, the DICOM software was modified to place C-MOVE requests in a separate database. Included in a database record are the level of the data requested (PATIENT, STUDY, SERIES, or IMAGE), the UID of the data, the destination, and the date and time of the request. A database entry is formulated and its uniqueness is verified (the criteria for duplication are level, UID and destination). Any duplicate requests generate a C-MOVE response to the requestor with a status of 'Cancel'. If a database record is written, the DICOM server returns a C-MOVE response to the requestor with a status of 'Success' and the number of failed operations set to zero.

A separate daemon monitors this database and fulfills all requests that are entered in it. Series of data are transferred by separate sub-processes which monitor the status of each transfer. Any failed transfers are reserved and repeated at a later time. This database C-MOVE facility is used by non-DICOM applications such as the PACS browser to make C-MOVE requests directly.

To increase the reliability and performance of the system, another application running from a different server records the status of all DICOM storage SCPs on the PACS network by checking the result of a DICOM echo to each server's designated node and port number. This database of status information is used by DICOM applications to bypass storage SCPs that are currently unavailable.

2.3.3. Image Conversion: DICOM to JPEG

A modality server can optionally perform 'teleradiology' functions which involves the conversion of DICOM images to JPEG format for viewing in a Web browser via resident WWW server software. Incoming images are tagged for conversion and a separate daemon processes the images as they are stored. The JPEG images are indexed and used by the WWW server application for in-line image display.

2.3.4. Plain Paper Printing

Each modality server and the retrieval server provide the facility to print on-line images on postscript laser printers located in various areas of the hospital. This provides a hard-copy alternative to soft-copy viewing for clinicians and patients. The server runs a printing daemon and a script that monitors print queues and gives feedback to users on the status of their print requests.

The retrieval server also has the facility to print images directly from CD should the data only be available in long-term storage.

2.3.5. Image Compression: Immediate or Delayed

Image data are compressed in short term-storage and on CD-ROM to both increase the amount of data on-line and reduce the number of CDs required for archival. As few of our DICOM entities support the DICOM JPEG compression standard, we apply lossless Lempel-Ziv coding (LZ77) to stored DICOM images using the UNIX program 'gzip'. This requires that any applications that deal with DICOM data include a facility to uncompress images to temporary DICOM files which can then be processed normally. The trade-off of processing and complexity for storage efficiency has halved the storage requirements for MRI and CT and cut storage requirements for ultrasound by 70 percent.

Image compression is delayed on modality servers and the retrieval server to move the high processing overhead of compression away from peak usage periods. Typically, the arrival of an image at the modality server triggers several processes that require access to the image. The processing overhead associated with these actions is reduced by leaving the images uncompressed until a later time. The cost of this performance enhancement is loss of storage space. A typical modality server, such the one serving the helical CT, stores about 2400 512x512 16-bit images each day which translates to about 1,200 MB of data on disk. With 2-to-1 compression (typical), this means that about 600 MB of extra storage is required to maintain the uncompressed data throughout the day. However, this represents less than 2 percent of the available storage on a 32 GB RAID-5 array. This is an acceptable sacrifice for a significant gain in performance.

3. DISCUSSION

3.1. Data Redundancy

One of the goals of our PACS project was to correct a weakness inherent in a conventional film system, the inevitable loss of patient data. To this end, we have attempted in our design to build as much data redundancy into the system as budget and network constraints will allow.

The use of RAID technology has simplified the process by making data recovery as easy as replacing a failed disk and rebuilding the array without losing any data. This level of protection applies both to image data and databases.

CD archives are protected by recording an exact copy of the CD data onto DAT. Lost or damaged CDs can be replaced by simply restoring that data from CD, re-imaging the directory structure and re-recording the image, a two hour process.

3.2. System Security

Security issues have been addressed in several ways:

- The DICOM-compliant systems use the *Association* class to authenticate each SCU. This prevents unauthorized hosts from accessing scanners or servers directly via the DICOM protocol.
- The mSQL databases employ host- and user-level access control to database servers to prevent unauthorized reading or writing of database records.
- The WWW server uses both the host-level security built into the httpd daemon and username/password (encrypted) authentication written into the WWW applications.

The WWW server employs a system that times out a user login session after a specified period of inactivity. This prevents simultaneous logins to the same account. As well, users are forced to periodically change their passwords to enhance WWW system security. It is especially important that that WWW server security be robust as this application is available through the Internet.

3.3. Benefits of In-House Development

There are several advantages to developing a PACS in-house:

- **Cost:** the overall cost of an in-house solution including salaries and capital expenditures is still a fraction of what vendors charge for a commercial solution.
- **Flexibility:** in-house development allows the development team to respond quickly to user needs and to tailor solutions to user requirements.
- **Ease of Integration:** having a development team on-site reduces the problems associated with integrating commercial hardware and software into the system. Vendor DICOM implementations are not always completely compatible and, as in our case, often require in-depth troubleshooting to identify points of failure.

Well into the PACS project we received a request from another institution to have certain scans automatically routed to their workstation. Our routing operates at the image level, distributing images to various destinations using a 'round-robin' method. This ensures that no single destination will slow the pace of image distribution to all workstations. It was discovered that the DICOM storage SCU implementation for this particular workstation could not deal with a series of images spread over multiple associations, rendering the received data useless. Since we had complete control of development and implementation on the transmitting end, we were able to develop a gateway machine (using a simplified version of the modality server software) that collated series before sending them to the workstation in question.

Of course, the in-house approach requires that a business plan be formulated and a PACS development team put in place. These are not easy tasks to accomplish and require determination on the part of the radiology staff and administration to see the project through.

4. CONCLUSION

Films were eliminated entirely in CT and MRI in February of 1997 and ultrasound has been filmless for more than two years. In addition, the ultrasound mini-PACS has been converted to DICOM conformance while maintaining backward compatibility with archived data.

The system continues to evolve as new technologies become available. As well, system performance is being examined by gathering usage and data transfer statistics to aid in the refinement of the PACS design and implementation².

By employing a distributed PACS design, we have spread network traffic and computational burdens evenly over inexpensive hardware in a comprehensive manner. The resulting system is a model for low-cost radiological image management.

ABSTRACT

This paper describes the acquisition of usage statistics and analysis of data-flow patterns for our DICOM-conformant PACS during its first year of operation. The system currently supports one MR, nine ultrasound, and three CT scanners, and has allowed our department to fully eliminate the production and use of film in these modalities. The aim was to not only aid in trouble-shooting, but to provide a means of examining usage patterns and quantifying data-flow and storage requirements to help identify where refinements should be focused. Of particular interest were statistics quantifying turn-around times for user-initiated transfers and retrieval from long-term storage; the quantities of data acquired, moved, retrieved and prefetched, (sorted by modality and anatomy); usage patterns of clinicians within the hospital, and the quantity of data accessed via a WWW interface to our PACS. The results have been instrumental in refining our physical network plan, modifying retrieval and transmission algorithms, and providing objective measures of the performance of our PACS. During the evolution of the system, the same data have allowed us to retrospectively examine whether certain modifications yielded improvements that were significant and whether expectations were met. The logging process continues since it is now relied upon as a tool for monitoring nearly all useful system parameters.

Keywords: PACS, usage statistics, performance, filmless radiology, DICOM

1. INTRODUCTION

The Montreal General Hospital (MGH) is a 500 bed level-1 trauma center and is the largest of four teaching hospitals that are affiliated with McGill University. During the last year, our department performed over 160,000 radiological examinations. Approximately 40 percent of these were in CT, MRI and ultrasound, while the rest comprised conventional x-ray, mammography, fluoroscopy and angiography.

In late 1995, we developed and implemented an ultrasound mini-picture archive and communications system (PACS) to improve access to and management of image data, and to eliminate film. Based on the success of this system, along with some strong clinical motivations, it was decided that it should be expanded to include CT and MRI. The PACS was subsequently redesigned to handle the larger volume of data and to conform to the Digital Imaging and Communications in Medicine (DICOM) standard. In January of 1997, the new system went into operation, supporting one MR, nine ultrasound, and two CT scanners. (A third CT scanner was added in May.) In February (1997), films in CT and MRI were eliminated.

Because the system was developed in-house, we had the freedom to incorporate extensive logging to trace the flow of data and monitor system usage. While the resulting information proved useful in trouble-shooting early on, it became clear that it had value in measuring performance parameters that relate to system design and usage. Thus, the logging facility has become an integral part of our PACS and presently records all significant user-initiated actions along with system-generated operations that result in the transmission of image data throughout the network.

This paper describes the acquisition of these data and the results covering the past year of operation. This information has provided insights that have helped guide the ongoing development of our PACS which, as described

Other author information: (Send correspondence to C.J.H.):

Dept. of Diagnostic Radiology
 Montreal General Hospital
 1650 Cedar Avenue
 Montreal, Quebec H3G 1A4
 Canada
 E-mail: chenri@rad.mgh.mcgill.ca



below is best characterized as a distributed model. Here we share the lessons learned through our design modifications along with figures pertaining to system usage, data volumes handled, and transmission rates. These figures should be useful in avoiding design pitfalls and helping to anticipate transmission and usage loads where models similar to ours are employed.

2. METHODOLOGY

2.1. Background

Our current DICOM PACS grew out of the experience gained in developing our own ultrasound mini-PACS¹ which was based predominately on a network of Macintosh computers. The latter system, which was not DICOM-conformant until recently, employs 'capture stations' equipped with framegrabbers next to each ultrasound scanner (interfaced via an NTSC video output on each scanner) which transmit the acquired images to a central 'server' where they become available for viewing from any Macintosh on the departmental network. A master database keeps track of every study performed and allows users to quickly locate and retrieve prior examinations. For long-term archival, we employ recordable compact disks.

While operating this mini-PACS, before deciding to include CT and MRI, several enhancements were introduced which ultimately are required by any effective PACS. These included fully automating the archival process (on CD) providing database searching tools that are integrated with the viewing software; providing a means of submitting and servicing requests for retrieval of data from long-term archive, and enabling access to the data from non-Macintosh computers outside the department. Together, these components satisfied the varied needs of the users and facilitated the management of image data to the point where film was eliminated; one film processor was closed, and one clerk (previously dedicated to film-management) was reallocated to duties elsewhere.

Concurrent to these developments, a group of our radiologists were exploring the merits of soft-copy review practices in CT and MRI. After a six month trial, they concluded empirically that film-based review was inferior to soft-copy analysis. Thus, together with the success of the US mini-PACS, the decision to pursue development of a CT and MRI PACS seemed well justified. This time, it was decided that the new system should adhere to the DICOM standard to facilitate future upgrades and ease the integration of imaging equipment that would also be DICOM conformant. This necessitated a complete redesign. (See our companion paper in these proceedings for more details.²)

2.2. PACS Design Overview

Apart from ten commercial dual-monitor diagnostic review workstations (all Unix-based), our entire PACS was developed in-house. Each component of the system is based upon IBM PC-compatible computers running the *LINUX* operating system (Red Hat distribution version 4.2). The choice to employ this combination of hardware and operating system stemmed partly from financial considerations (i.e., PCs are relatively low in cost, and *LINUX* is free), but the choice was also based upon performance criteria that precluded the use of Macintosh and Windows 95/NT systems.

The core of our DICOM software was derived from the Mallinckrodt Institute of Radiology Central Test Node software, which is freely available in the public domain. For database software, we used mini-SQL³, which is also available in the public domain. Several modifications were required to accommodate the archival of images on CD. In particular, a 'Master Database' was introduced along with several disk space management applications. Additional modifications were made to incorporate automatic routing and prefetching capabilities, and to make DICOM query/retrieve operations more robust and manageable by the CPU under heavy loads.

The PACS itself comprises several 'servers' dedicated to specific tasks. The functional roles of each server are summarized in Table 1.

The network presently runs 100MBits/sec Ethernet with every server and each workstation on its own port of an Ethernet switch. When first deployed, the PACS was based on 10MBits/sec Ethernet with 2-3 workstations sharing each port of a 10MBits/sec Ethernet switch.

The PACS presently serves 9 ultrasound scanners from various manufacturers, 3 CT scanners (2 of which required a commercial solution to become users of the DICOM Storage service class; the other is newer and has it built-in).

¹ Hughes Technologies, Pty., Ltd., Australia

Table 1. Summary of departmental PACS servers and their functional roles.

Modality Server	<ul style="list-style-type: none"> - 1 for each scanner - DICOM Storage SCU/SCP - Query/Retrieve SCP 	<ul style="list-style-type: none"> - accepts images from scanner - provides short-term storage - auto-routes images to appropriate destinations - auto-prefetches - services query/retrieves - accepts images from workstations - services print requests (paper) - provides WWW image viewing
Archival Server	<ul style="list-style-type: none"> - 1 for department - DICOM Storage SCP 	<ul style="list-style-type: none"> - accepts images from Modality Servers - automatically records data on CDs
Retrieval Server	<ul style="list-style-type: none"> - 1 for department - DICOM Storage SCU 	<ul style="list-style-type: none"> - restores images from CD for re-distribution to other DICOM destinations - services print requests (paper) - provides WWW image viewing - coordinates all manual prefetching and scheduled routing
Database Server	<ul style="list-style-type: none"> - 1 for department - no DICOM functions 	<ul style="list-style-type: none"> - maintains Master Database to keep track of all studies performed - knows location of any study/series eg, on-line, off-line (CD number)
WWW Server	<ul style="list-style-type: none"> - 1 for department - DICOM Storage SCP 	<ul style="list-style-type: none"> - provides WWW access to all images and Master Database - coordinates viewing image series, downloading, moving, searching, printing - several tools available

and 1 MR scanner (a DICOM Storage SCU). Modality Servers are distributed as follows: One for MRI, one for the two older CT scanners; one for the third CT scanner, which is located in the Emergency Department, and one that serves the 9 ultrasound scanners. Images are transmitted automatically by every CT and MR scanner (and every US capture station) to the corresponding Modality Server where they are stored for several weeks (approx. 90 days). As each image is received, a copy is automatically routed to one or more DICOM destinations according to pre-defined rules that, in our case, depend primarily on the imaged anatomy. Since our radiologists are specialized by anatomy (rather than by modality, for example), we have chosen to designate the diagnostic review workstations accordingly. Thus, we have 2 abdominal; 2 bone/spine; 1 chest; 1 ENT and 2 neurological workstations, plus one dedicated to MRI and one dedicated to emergency CT cases. A copy of every scanned image is transmitted to the Archival Server to be recorded on CD, and until recently, a copy of every image was transmitted to the WWW Server to be available for viewing through the Internet. WWW access to the images is now distributed among the Modality Servers, allowing more images to remain immediately accessible and eliminating the need to transmit every image to the WWW Server.

Each Modality Server provides several weeks of storage on hard disks that recently have been converted to employ RAID technology (level 5). After six months of operation, we decided to introduce the use of lossless compression

to extend the lifetime of the data on-line (i.e., on each Modality Server) as well as increase the effective capacity of the CDs used for long-term archival. The data, sorted by modality, are recorded on CD typically within 24 hours except for ultrasound which takes three to four days to fill one CD¹. Once recorded and labeled, the CDs reside in the filing room in close proximity to the Retrieval Server which is equipped with two high-speed CD-ROMs.

From the beginning, it was decided that we would examine the need for juke-box technology that would allow data retrieval to be performed automatically. Thus, it was important to gather figures concerning the number of CDs handled each day and the times required to manually load them following audible notification.

The WWW Server is a key component of our PACS.³ Before ceasing to print films in CT and MRI, we recognized the need to provide a means of allowing clinicians throughout our institution, as well as referring physicians from outside, to access the images. The decision to employ the Internet stemmed from several factors, including the growing familiarity of the population with the World Wide Web; the relatively intuitive interfaces of Web-browsers; the readily available infrastructure; the desire to provide platform independent accessibility; and the ease of developing WWW interfaces. Cost was also a factor. By employing a Web-based interface to our PACS, we immediately made it possible to access images from any networked PC in the entire hospital without the additional costs associated with commercial viewing software.

While the WWW Server was originally intended to serve the image viewing needs of users outside our department, it quickly grew to provide several PACS management tools that are now used by almost everyone. These include the ability to: move any study or series for any patient to any DICOM device known to our system (eg, workstations); download any series in its full resolution for analysis off-line; manually schedule the prefetching of prior studies; print any series (on paper only); and create teaching-files. Several other Web-based tools are employed to administer user accounts; modify automatic routing rules; monitor system usage and examine database contents.

The Database Server maintains a 'Master Database' to record every patient examined and every study and series performed. It also maintains a record of the location(s) where the data reside (for example, which Modality Server they are on, and which CD they are stored on), and whether the data are on- or off-line.

2.3. Data-Flow and Usage Measurements

At the time of writing, we are completing the integration of our PACS and hospital information system (HIS). Once put into operation, the link between images and reports will be available (through the WWW Server), along with additional related information. We expect system usage figures to increase dramatically, but we are presently unable to speculate on how much.

2.3.1. Data acquisition

The data studied in this paper can be grouped into one of three categories: Those related to data transfer rates and volumes; those describing retrieval response times; and those related to the frequency and type of user interactions (through the WWW Server). In every case, the data were gathered by making use of the logging facilities that are embedded in almost every application comprising the PACS. The mini-SQL database software was used for this purpose, greatly facilitating the retrieval of the desired quantities.

2.4. Data Volumes

Among the first questions we sought to answer was whether the departmental diagnostic review workstations had been allocated equitably among the various anatomical specialties. This was a simple matter of examining data volumes in terms of the number of images acquired for each specialty. Although these numbers could be obtained from any one of several sources, they were most easily obtained by consulting the Master Database.

Next, we were interested in obtaining figures describing the amount of image-based network traffic. This information, we believed, would be useful in balancing the load across the network by identifying whether a given Modality Server was over- or under-utilized. In this case, we examined the number of user-initiated 'moves' (of series) from each Modality Server. Note that these 'moves' could be accomplished either through the WWW Server or through query/retrieval operations. The required data were obtained by recording details like the size, destination, and start and end transmission times for every series moved to a workstation or other DICOM device.

¹ Before employing compression, we recorded one CD per day for each Modality Server, including ultrasound.

2.5. Data Transfer Rates

We were also interested to know more about image transfer rates between our servers and the workstations. Here, we examined only the data transfer rates for user-initiated 'moves'. We recorded the time when the user issued the 'move' command for a given series (i.e., start-time), and the time when the entire series had been transmitted (i.e., end-time). All times were measured to the nearest second. Transfer rates were computed from the size of each image transmitted and the difference between start- and end-times. As discussed below, these measurements were performed before and after some significant modifications were made to the network hardware and transmission software.

2.6. Retrieval From Archives

In order to study retrieval patterns from long-term storage, we gathered figures quantifying the number of CDs handled; the times to respond to retrieval requests (measured from the time a request is received to the time a CD is manually loaded), and the frequency of requests for data versus time since the study was performed. The first quantity was intended to provide an indication of the workload imposed upon the filing room staff, while the second quantity measured their performance. Although our PACS had been in operation for only one year, the third study was intended to illustrate the degree to which older data are required which, in turn, has implications concerning the amounts of on-line storage provided by the Modality Servers.

2.7. WWW Server Usage

In terms of usage measurements, we were limited to the actions that are recorded for every user of the WWW Server. These include tracking the identity of the user logged-in, along with the IP address of their computer; what database queries are performed; what studies and series are selected; what series are viewed, moved, printed, downloaded, prefetched, or used to create a teaching file, and whether the desired data were available or had to be retrieved from CD. Every action is recorded with a time and date stamp to facilitate the generation of audit trails.

As simple measures of overall use, we extracted several summary statistics over the most recent eight months. These include the monthly totals of: new user accounts; logins including at least one query; series viewed; on- and off-line series moved (from a Modality Server to any known DICOM device); series printed, and series prefetched.

The WWW Server permits users to view any series of images whether on- or off-line. In the latter case, once the user asks to view a series, a request is sent immediately to the Retrieval Server and the user must wait for the data to be retrieved from CD (i.e., copied to the disks on the Retrieval Server). The images are then able to be viewed in JPEG format or downloaded in their native raw DICOM format. Images residing on the Modality Servers are available in both DICOM format and JPEG (the latter are immediately available for viewing). The generation of the JPEG images requires a conversion process from their original DICOM format. This process occurs automatically when each DICOM image is received on a Modality Server, or restored from CD. Thus, the selection of window widths and levels is automated. In many cases, we had to derive suitable defaults for the various combinations of anatomy imaged and modality. Some minor image processing is also employed. When viewed, the user is able to specify their own window width and level, as well as resize and sort the images. Changes in the window width or level require that the WWW Server regenerate temporary versions of the JPEG images from the corresponding full resolution DICOM images.

As an indicator of whether the default window widths and levels are satisfactory, we examined the number of times users chose to make changes themselves. Because most users may not succeed in selecting appropriate values on their first try, we counted only one occurrence per displayed series per login session.

Finally, we were interested in obtaining an appreciation for the degree to which users actively use their accounts on the WWW Server. For this purpose, we defined an account as 'active' if, at any given time, it had been accessed on at least two separate occasions within the most recent two weeks. In order to assess any trends, we determined the number of users with active accounts over the past eight months by taking measurements once per week.

3. RESULTS

At the time of writing, our PACS archive holds over 72,000 studies from 47,018 patients stored on approximately 1500 CDs. The total number of registered users authorized to access our WWW Server is 481; 431 of which are affiliated with the McGill University Health Centre, while the rest (50) are primarily outside referring physicians.

3.1. Data Volumes

The data collected to study the issues described above are presented in the tables below. Figures describing the volumes of image data generated in CT and MRI for each anatomical specialty are presented in Table 2. Equivalent figures for ultrasound are omitted since the examinations, in our department, are typically not characterized by anatomical specialty. The figures for each anatomy were collected over a three month period and include the mean number of studies performed per week; the total number of images per week and the mean number of images per study.

Table 2. Volumes of data generated per week per anatomical specialty. Percentages of the totals in each modality are expressed in parentheses.

	Abdomen	Bone	Chest	ENT	Neuro	Spine	Other
# studies (CT)	69 (21%)	12 (4%)	43 (13%)	37 (12%)	118 (36%)	39 (12%)	7 (2%)
# images (CT)	6297 (24%)	1735 (7%)	2894 (11%)	5441 (21%)	4476 (17%)	4689 (18%)	190 (1%)
mean images/study	91	142	68	145	38	122	27
# studies (MR)	15 (12%)	22 (17%)	0 (0%)	8 (6%)	45 (35%)	21 (16%)	18 (14%)
# images (MR)	2610 (17%)	2473 (16%)	0 (0%)	684 (4%)	6526 (41%)	2345 (15%)	1159 (7%)
mean images/study	174	110	0	91	144	110	66
Totals							
#studies (CT+MR)	84 (19%)	34 (7%)	43 (9%)	45 (10%)	163 (36%)	60 (13%)	25 (6%)
#images (CT+MR)	8907 (21%)	4208 (10%)	2894 (7%)	6125 (15%)	11002 (26%)	7034 (17%)	1349 (3%)

Table 3 presents the mean number of series transmitted from each CT and MR Modality Server, per day, as a result of user-initiated requests. The table includes results for both before and after the introduction of compression. Note that move requests may be generated either through the WWW Server or via direct DICOM query/retrieve operations.

Table 3. Volumes of data transmitted from each Modality Server as a result of user-initiated 'moves' (i.e., via the WWW Server or through DICOM query/retrieves). 'BC' = Before compression. AC = After compression.

	MR Server	CT Server #1	CT Server #2
mean # of series moved (BC)	144/day	70/day	70/day
mean # of series moved (AC)	48/day	43/day	42/day
median series size	2.2 MB	7.8 MB	11 MB

3.2. Data Transfer Rates

The transfer rates for user-initiated moves of images (or series) are presented in Table 4. Results are shown before and after some significant modifications were made to both the network and the software used to transmit the image data.

3.3. Retrieval From Archives

The results pertaining to the retrieval of data from long term archive are as follows. The number of CDs requiring manual loading was determined by computing a daily average over the most recent one month period. Only weekdays were included. A total of 69 CDs (average) were loaded per day; 28 of which were for CT data; 12 for MRI, and 29 for ultrasound. Prior to compressing data on both CD and on the Modality Servers, the figures for CT and MRI were 46 and 19, respectively. The figure given for ultrasound (29), in fact, does not yet reflect the effect of compression since it was introduced only recently. In other words, we expect to see it decrease.

Table 4. Image data transfer rates (in kilobytes/second) between Modality Servers and workstations for user-initiated 'moves'. Results are shown for transfer rates before and after modifications that included converting from 10 to 100Mbits/sec Ethernet; changing the physical network layout, and making software enhancements. Note: The results for the MR Server 'before modifications' were obtained while it employed 10Mbits/sec Ethernet. All others employed 100Mbits/sec.

	MR Server	CT Server #1	CT Server #2
Before Modifications			
median kbytes/sec	32	88	83
mean kbytes/sec	39	91	88
After Modifications			
median kbytes/sec	205	201	190
mean kbytes/sec	239	238	209

The median time elapsed between receiving a retrieval request and loading the required CD was 349 seconds for requests submitted between 7-17:00 hours. Between 17-24:00 hours, the median was 2,250 seconds. (The distinction between the day-shift and night-shift is explained in Section 4 below.) From midnight until 7AM, our filing room is closed and retrieval requests are not serviced. These figures were determined by monitoring all such activity over the most recent month.

Figure 1 presents a histogram of the number of requests for data as a function of the time between when the examination was performed and when the request was submitted. These data were acquired as follows: For every series moved, printed or prefetched by a user through the WWW Server, we determined the number of hours between the date and time the user issued the request and the date and time the series was created (i.e., when the patient was imaged). The data were then binned in 24 hour intervals to generate the histogram. The thin curve at the top of the plot shows the integral sum of the number of series needed to date.




Figure 1. Histogram plot illustrating the number of series required at any given time as a function of the time since the series were acquired. The thin curve at the top of the plot shows the integral sum of the number of series needed to date.

3.4. WWW Server Usage

The monthly totals of new user accounts created are listed in Table 5. The number of logins (plus at least one database query) per month are listed in Table 6. Also listed is the number of distinct users who logged-in and made at least one query, and the median number of logins per user per month. Table 7 lists the monthly total number of series viewed through the WWW Server, along with the corresponding number of distinct users and the median number of series viewed per user per month. Table 8 presents the equivalent figures for the total number of on-line and off-line series moved (eg. from a Modality Server to a workstation). The number of series printed (on paper) per month, and the monthly totals of series prefetched are listed in Tables 9 and 10, respectively. (Note: Printing on paper through the WWW Server became available in July. Prefetching became available in June.)

Table 5. Monthly total number new user accounts created.

May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
28	38	50	41	48	38	37	29

Table 6. Monthly total number of logins (including at least 1 database query).

	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
total logins	907	837	2792	4120	5371	5765	5897	5733
# distinct users	72	64	115	134	163	168	180	181
logins/month/user (median)	5	3	10	14	14	18	16	11

Table 7. Monthly total number of series viewed through the WWW Server.

	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
total series viewed	2155	1621	2335	1954	2630	4248	4507	3719
# distinct users	93	90	98	103	134	132	147	141
series/month/user (median)	9	6	6	7	7	12	15	9

The number of times users chose to make changes to the default window widths and levels that are employed when viewing a series of images through the WWW Server was examined over the most recent 4 months. A total of 15,104 series were viewed. Users chose to modify the window widths and levels at least once in 234 series.

4. DISCUSSION

It is important to realize when examining the results that there is a large number of influencing factors. Some explanations are simple, or obvious, like correlating holiday periods with system usage, or changing technical parameters like CPU or network speeds. Others are much more difficult to identify. It is difficult, for example, to determine whether the increasing use of the WWW Server is due to growing user acceptance or simply the growing number of users. Another confounding factor is that the PACS is continuing to evolve, providing more tools and better interfaces to existing ones. So rather than attempt to isolate some of the more complex causes and effects, we have focussed on what we believe to be larger details.

The first among these was to verify that the departmental diagnostic review workstations had been allocated equitably. Using Table 2 to rank the various anatomical specialties on the basis of total studies (CT+MR) yields the following (from highest to lowest): Neuro (163); Bone/Spine (94); Abdomen (84); ENT (45); Chest (43) and Other

Table 8. Monthly total number of on-line and off-line series moved.

	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
on-line series moved	1097	950	3531	4274	7506	5986	7335	5617
# distinct users	53	67	75	86	113	107	118	116
series/month/user (median)	6	6	11	14	8	11	11	8
off-line series moved	2930	1715	2808	3672	5037	5949	6180	3726
# distinct users	63	71	72	76	101	114	119	95
series/month/user (median)	10	7	7	10	12	17	11	10

Table 9. Monthly total number of series printed (on paper).

	Jul	Aug	Sep	Oct	Nov	Dec
total series printed	398	814	900	886	939	735
# distinct users	19	34	37	47	53	48
series/month/user (median)	3	4	5	6	7	6

(25). Neuro and Bone/Spine exchange their rankings if the total number of images is considered. Given that each specialty must be provided with at least one workstation and that there must be one dedicated workstation for MRI and one for the Emergency CT scanner, 3 are left to be distributed. Giving one each to Abdomen, Bone/Spine and Neuro is most equitable, which confirms that the workstations were allocated properly to begin with.

During the first four months of operating our PACS with CT and MRI, we had only two Modality Servers: One shared by our CT scanners, and one for MRI. Upon examining Table 3, which describes the number of series moved by users from each Modality Server on a daily basis, it appears at first that the two CT Modality Servers perform half the work of the MR Server. However, by taking into account the size of the series transferred, it is clear that the CT Servers move a greater volume of data (measured in MB) each day than is moved in MRI.

After introducing compression to increase the amount of data residing on-line, the number of series moved decreased by almost a factor of two in CT, and a factor of three in MRI. At first, this might seem unexpected since having more data on-line should increase the likelihood of a series being moved. The explanation, we believe, is two-fold: First, the local storage on the WWW Server was increased by adding the use of compression along with one 9GB hard disk. At the time, images could only be viewed from the WWW Server. So they had to be transmitted there if they were not already present. Thus, the larger capacity of the WWW Server reduced the number of images needed to be transferred from the Modality Servers. Secondly, if the previous study for a patient undergoing a repeat examination was still on-line and was required for comparison, the manual prefetching command generated the equivalent of a 'move'. With the addition of automatic prefetching (coinciding with the introduction of compression), fewer studies had to be moved manually. While the same quantity of data might still flow from the Modality Servers, the relative volume of data moved manually has decreased.

The results presented in Table 4 showing the image data transfer rates for user-initiated moves are particularly

Table 10. Monthly total number of series prefetched.

	Jun	Jul	Aug	Sep	Oct	Nov	Dec
total series prefetched	304	671	798	956	1347	1326	2288
# distinct users	3	4	4	3	3	4	9

interesting. In the 'before modifications' category, the results for the MR Modality Server are surprisingly poor. One explanation is that these results were gathered while the server was still on a 10MBits/sec Ethernet segment shared by the MRI scanner, and the two were linked to the rest of the PACS by one coaxial cable that was shared by several other computers (non-servers) in the department. The corresponding results for the two CT Modality Servers were obtained after they had been upgraded to 100MBits/sec Ethernet, but the relative improvement over the MR Server was only a factor of approximately 3. Still, the results were surprisingly poor (approx. 90kb/sec). The explanation for these results, and the more than two-fold increase shown in the 'after modifications' category, lies in a change that was made to the application responsible for transmitting the images. Previously, this application had been over-engineered in an attempt to reduce the load on the Modality Server CPUs and to provide complete immunity to transmission interruptions. These goals were achieved by keeping a record of which images had been transmitted successfully and permitting only a single image to be transmitted at a time. This meant that a DICOM association was opened and closed for every image sent. The resulting overhead is most significant when transmitting large series of images and accounts for the factor of two increase in speed. The apparent six-fold increase for MRI is due to both the change from 10MBits/sec Ethernet and the newer algorithm.

The figures describing the number of CDs manually loaded per day illustrate vividly the benefits of compression. The argument for not employing compression earlier had been that it would slow down the already slow process of restoring data from CD to the local hard disk. After conducting some tests to determine the additional fraction of time required to uncompress several large series of images, we were surprised to find that there was no measurable difference. However, the time required to compress the images was much greater than expected, placing a significant burden on the CPU if performed during peak hours.

The median time required to load a CD during the day (349 seconds) was higher than expected. Beyond 17:00 hours, only one clerk remains on duty and prefetching of previous ultrasound cases is done for the following day. (Prefetching in CT and MRI is fully automated.) The dramatic increase in response time (2250 seconds = 37.5 minutes) was also surprising and seems to be due to a reduced sense of urgency to service the requests. Consequently, the prefetching task has been spread more evenly across the day and night shifts.

The histogram in Figure 1 illustrates the frequency of requests for images as a function of the difference in time between the corresponding study date and when the request was issued. The plot is intended to provide an indication of the period of time most data are needed following a given examination. Knowing the length of time images remain on-line on the Modality Servers, this plot should allow the volume of off-line data requested to be predicted. Conversely, if a maximum number of studies to be retrieved from CD per day is specified, the plot should allow the optimum amount of disk space allocated for on-line storage to be determined. It can be seen that the current storage capacity of our Modality Servers (approx. 90 days) is sufficient to account for nearly 75 percent of the data required. Since we would like this figure to approach to 90 percent, we would have to double the storage capacity of the Modality Servers.

The results collected describing the usage of the WWW Server were also very telling. With nearly 500 registered users, the total continues to grow at over 30 new accounts per month (not counting December because of holidays). The rise in the number of new user accounts created during the summer months was easily explained: The new clinical residents arrived in June, and 15 new PC-based workstations were provided for image viewing in several of the clinics and operating theatres in the hospital.

Table 6 shows the increasing number of people using the system (i.e., logging-in and performing at least one database query). The most rapid growth occurred during the summer, going from 907 logins per month in May to over 5000 per month by September. The number of distinct users per month and their median number of logins grew in the same manner. By the end of the summer, the number of distinct users per month was approximately 45 percent of the total number of registered users, and each employed the WWW Server (on average) every other day.

A similar trend is seen in Table 7 which shows the number of image series viewed per month. Table 8 shows the number of on-line and off-line series moved per month by users. Although the numbers seem to reach a plateau, they are alarmingly high given that they do not include network traffic due to automatic image routing and prefetching. It turns out that a large number of these moves were directing images to the WWW Server to be viewed. Thus, in order to reduce (in fact, eliminate) this traffic, we decided to distribute the WWW image display task among the Modality Servers and the Retrieval Server where most of the images reside. Although we do not have results, we expect the figures describing the number of series moved to change dramatically. This modification may also affect




Figure 2. Bar plot illustrating the number of 'active' users of the WWW Server during the weeks from May through December.

several other results. The reduced network traffic may improve image transfer times, and the immediate availability of images for viewing where previously the user had to wait may improve user acceptance.

Table 9 shows the monthly total number of series printed on paper. The rapid jump from 398 to 814 series between July and August coincided with the introduction of several new laser printers (dedicated to printing images) in key locations throughout the hospital. The current number of series printed per month (approx. 900) is disappointing. This number is approximately 25 percent of the number of series viewed through the WWW Server. The most likely reasons for these results are that a significant number of users have not yet become accustomed to soft-copy viewing (or are unwilling to try); they lack access to an adequate computer, or they are dissatisfied with the display function itself. Certainly, the inability to modify the displayed window width and level settings in real-time is a barrier. In order to address this issue, we will soon be distributing more sophisticated image viewing software throughout our hospital.

The number of series prefetched per month and the number of distinct users doing the prefetching are presented in Table 10. From June to October there is a steady increase in the number of series prefetched, then it appears to stabilize in November. In December, the figure suddenly doubles. The latter is easily explained: The ultrasound mini-PACS was integrated with the CT and MR in December, and both now employ the same prefetching scheme. Thus, the surplus of prefetched series is attributable to ultrasound. The increase in the number of distinct users doing prefetching reflects the participation of our ultrasound technicians who were accustomed to the older PACS. The steady increase noted from June to October is explained as follows: Image data in CT and MRI began to be archived on CD in January. Thus, although the same number of studies were performed every month, the amount of digital data that could be prefetched from CD had to accumulate. In the earliest days, for example, the vast majority of previous cases had existed only on film.

Out of 15,104 series viewed through the WWW Server over the past 4 months, the user changed the window width and levels 234 times (i.e., 2 percent of the time). This would seem to indicate that the default values were acceptable the great majority of the time. However, it might also indicate that most users do not understand how, or have the patience, to change the values manually.

Finally, the increase in use of the WWW Server over the months is clearly seen in Figure 2. While some of the other results are sensitive to the actions of single users, this quantity is a more accurate reflection of the increasing load due to new users and those that may use the WWW Server more often over time.




Figure 2. Bar plot illustrating the number of 'active' users of the WWW Server during the weeks from May through December.

several other results. The reduced network traffic may improve image transfer times, and the immediate availability of images for viewing where previously the user had to wait may improve user acceptance.

Table 9 shows the monthly total number of series printed on paper. The rapid jump from 398 to 814 series between July and August coincided with the introduction of several new laser printers (dedicated to printing images) in key locations throughout the hospital. The current number of series printed per month (approx. 900) is disappointing. This number is approximately 25 percent of the number of series viewed through the WWW Server. The most likely reasons for these results are that a significant number of users have not yet become accustomed to soft-copy viewing (or are unwilling to try); they lack access to an adequate computer, or they are dissatisfied with the display function itself. Certainly, the inability to modify the displayed window width and level settings in real-time is a barrier. In order to address this issue, we will soon be distributing more sophisticated image viewing software throughout our hospital.

The number of series prefetched per month and the number of distinct users doing the prefetching are presented in Table 10. From June to October there is a steady increase in the number of series prefetched, then it appears to stabilize in November. In December, the figure suddenly doubles. The latter is easily explained: The ultrasound mini-PACS was integrated with the CT and MR in December, and both now employ the same prefetching scheme. Thus, the surplus of prefetched series is attributable to ultrasound. The increase in the number of distinct users doing prefetching reflects the participation of our ultrasound technicians who were accustomed to the older PACS. The steady increase noted from June to October is explained as follows: Image data in CT and MRI began to be archived on CD in January. Thus, although the same number of studies were performed every month, the amount of digital data that could be prefetched from CD had to accumulate. In the earliest days, for example, the vast majority of previous cases had existed only on film.

Out of 15,104 series viewed through the WWW Server over the past 4 months, the user changed the window width and levels 234 times (i.e., 2 percent of the time). This would seem to indicate that the default values were acceptable the great majority of the time. However, it might also indicate that most users do not understand how, or have the patience, to change the values manually.

Finally, the increase in use of the WWW Server over the months is clearly seen in Figure 2. While some of the other results are sensitive to the actions of single users, this quantity is a more accurate reflection of the increasing load due to new users and those that may use the WWW Server more often over time.

5. CONCLUSIONS

This paper has presented an overview of our DICOM-conformant PACS along with an analysis of data-flow and usage statistics gathered during its first year of operation. The generation of these results was facilitated by having incorporated detailed logging of transmission events and user-initiated actions directly into the PACS and its WWW interface. Initially, the rationale for gathering such detailed information derived from the need to facilitate debugging and to collect audit trails for every user to satisfy hospital security concerns. It soon became clear that this information also allowed us to objectively assess the performance of the PACS while it continued to evolve.

In several cases, analysis of the information gathered allowed us to identify system faults and weaknesses that did not manifest themselves openly. Examples of these include discovering that one Modality Server was operating on a 10MBits/sec Ethernet segment when it was thought to be working at 100MBits/sec. An analysis of image transfer rates identified this problem. On another occasion, automatic image routing was accidentally confusing ENT and emergency cases such that the emergency cases were (erroneously) routed to the ENT workstation. An analysis of data volumes routed identified this error. At one point, an inordinately high number of CD retrievals were repeatedly issued over a period of several days. A clerk in a different department was later identified from the WWW Server usage logs as being responsible. The confusion in this case was due to a misunderstanding over which cases were needed to prepare for rounds.

The statistics we gather also allow us to anticipate the load placed on the PACS by the increasing total number of users. A recent design modification has distributed the WWW image display task across the Modality Servers. This new scheme minimally increases the CPU load on each Modality Server, but it virtually eliminates all network traffic (i.e., DICOM image transfers) due to users from outside of our department. It has also liberated the WWW Server to handle only WWW connections and database queries while previously it also received DICOM images from every Modality Server and converted them to JPEG format images. We are now gathering information to study the effects of this modification.

In this respect, building the logging facility into our PACS has been invaluable. It too continues to evolve not just providing information concerning new features of the PACS, but is being refined to capture more detailed information in ways that make analysis easier. The recently implemented gateway between our PACS and HIS will allow us to provide more than just images to our users, so we expect a significant increase in system usage. The ability to monitor this in the months ahead should prove useful in adapting the new utilities to the needs of the users and rapidly addressing any oversights.